Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4 Effective date of registration: 11/26/2014	*read-view* — (Optional) String of a maximum of 64 characters that is the name of the view. The default is that the read-view is assumed to be every object belonging to the Internet object identifier (OID) space (1.3.6.1), unless the **read** option is used to override this state. **write** — (Optional) Specifies a write view for the SNMP group. This view enables you to enter data and configure the contents of the agent. *write-view* — (Optional) String of a maximum of 64 characters that is the name of the view. The default is that nothing is defined for the write view (that is, the null OID). You must configure write access. **notify** — (Optional) Specifies a notify view for the SNMP group. This view enables you to specify a notify, inform, or trap. *notify-view* — (Optional) String of a maximum of 64 characters that is the name of the view. By default, nothing is defined for the notify view (that is, the null OID) until the **snmp-server host** command is configured. If a view is specified in the **snmp-server group** command, any notifications in that view that are generated will be sent to all users associated with the group (provided a SNMP server host configuration exists for the user). Cisco recommends that you let the software autogenerate the notify view. See the "Configuring Notify Views" section in this document. **access** — (Optional) Specifies a standard access control list (ACL) to associate with the group. **ipv6** — (Optional) Specifies an IPv6 named access list. If both IPv6 and IPv4 access lists are indicated, the IPv6 named access list must appear first in the list. *named-access-list* — (Optional) Name of the IPv6 access list. *acl-number* — (Optional) The *acl-number* argument is an integer from 1 to 99 that identifies a previously configured standard access list. Cisco IOS SNMP Support Command Reference (2013), at 343-44	Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1994

Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4 Effective date of registration: 11/26/2014	**snmp-server host** | Release	Modification
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.
15.2(1)S	This command was modified. The p2mp-traffic-eng notification-type keyword was added. | **Usage Guidelines** If you enter this command with no optional keywords, the default is to send all notification-type traps to the host. No informs will be sent to the host. The **no snmp-server host** command with no keywords disables traps, but not informs, to the host. To disable informs, use the **no snmp-server host informs** command. **Note** If a community string is not defined using the **snmp-server community** command prior to using this command, the default form of the **snmp-server community** command will automatically be inserted into the configuration. The password (community string) used for this automatic configuration of the **snmp-server community** command will be the same as that specified in the **snmp-server host** command. This automatic command insertion and use of passwords is the default behavior for Cisco IOS Release 12.0(3) and later releases. However, in Cisco IOS Release 12.2(33)SRE and later releases, you must manually configure the **snmp-server community** command. That is, the **snmp-server community** command will not be seen in the configuration. Cisco IOS SNMP Support Command Reference (2013), at 354	**snmp-server host** The snmp-server host command specifies the recipient of Simple Network Management Protocol (SNMP) notifications. Recipients are denoted by host location and community string. The command also specifies the type of SNMP notifications that are sent: a *trap* is an unsolicited notification; an *inform* is a trap that includes a request for a confirmation that the message is received. The configuration can contain multiple statements to the same host location with different community strings. For instance, a configuration can simultaneously contain all of the following: • snmp-server host host-1 version 2c comm-1 • snmp-server host host-1 informs version 2c comm-2 • snmp-server host host-1 version 2c comm-3 udp-port 666 • snmp-server host host-1 version 3 auth comm-3 The no snmp-server host and default snmp-server host commands remove the specified host by deleting the corresponding snmp-server host statement from the configuration. When removing a statement, the host (address and port) and community string must be specified. Platform all Command Mode Global Configuration Command Syntax snmp-server host *host_id* [*VRF_INST*] [*MESSAGE*] [*VERSION*] *comm_str* [*PORT*] no snmp-server host *host_id* [*VRF_INST*] [*MESSAGE*] [*VERSION*] *comm_str* [*PORT*] default snmp-server host *host_id* [*VRF_INST*] [*MESSAGE*] [*VERSION*] *comm_str* [*PORT*] Parameters • *host_id* hostname or IP address of the targeted recipient. • *VRF_INST* specifies the VRF instance being modified. — <no parameter> changes are made to the default VRF. — vrf *vrf_name* changes are made to the specified user-defined VRF. • *MESSAGE* message type that is sent to the host. — <no parameter> sends SNMP traps to host (default). — informs sends SNMP informs to host. — traps sends SNMP traps to host. • *VERSION* SNMP version. Options include: — <no parameter> SNMPv2c (default). — version 1 SNMPv1; option not available with informs. — version 2c SNMPv2c. — version 3 noauth SNMPv3; enables user-name match authentication. — version 3 auth SNMPv3; enables MD5 and SHA packet authentication. — version 3 priv SNMPv3. HMAC-MD5 or HMAC-SHA authentication. AES or DES encryption. • *comm_str* community string (used as password) sent with the notification operation. Although this string can be set with the snmp-server host command, the preferred method is defining it with the snmp-server community command prior to using this command. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1995

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4

Effective date of registration: 11/26/2014 | SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when it receives traps. The sender cannot determine if the traps were received. However, an SNMP entity that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the sender never receives the response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination than traps.

Compared to traps, informs consume more resources in the agent and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once; an inform may be tried several times. The retries increase traffic and contribute to a higher overhead on the network.

Cisco IOS SNMP Support Command Reference (2013), at 354 | SNMP notifications are messages, sent by the agent, to inform managers of an event or a network condition. A *trap* is an unsolicited notification. An *inform* (or inform request) is a trap that includes a request for a confirmation that the message is received. Events that a notification can indicate include improper user authentication, restart, and connection losses.

Traps are less reliable than informs because the receiver does not send any acknowledgment. However, traps are often preferred because informs consume more switch and network resources. A trap is sent only once and is discarded as soon as it is sent. An inform request remains in memory until a response is received or the request times out. An inform may be retried several times, increasing traffic and contributing to higher network overhead.

Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1963 |
| Cisco IOS 15.4

Effective date of registration: 11/26/2014 | **snmp-server source-interface**

To specify the interface from which a Simple Network Management Protocol (SNMP) trap originates the informs or traps, use the snmp server source interface command in global configuration mode. To remove the source designation, use the no form of this command.

snmp-server source-interface {traps | informs} *interface*
no snmp-server source-interface {traps | informs} [*interface*]

Cisco IOS SNMP Support Command Reference (2013), at 376 | **snmp-server source-interface**

The snmp-server source-interface command specifies the interface from which a Simple Network Management Protocol (SNMP) trap originates the informs or traps.

The no snmp-server source-interface and default snmp-server source-interface commands remove the inform or trap source assignment by removing the snmp-server source-interface command from running-config.

Platform   all
Command Mode   Global Configuration

Command Syntax
snmp-server source-interface INTERFACE
no snmp-server source-interface
default snmp-server source-interface

Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1967 |

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4 Effective date of registration: 11/26/2014	 Cisco IOS SNMP Support Command Reference (2013), at 394	 Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1999

Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4 Effective date of registration: 11/26/2014	**Usage Guidelines** To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides. Also, before you configure remote users for a particular agent, configure the SNMP engine ID, using the snmp-server engineID command with the remote keyword. The remote agent's SNMP engine ID is needed when computing the authentication and privacy digests from the password. If the remote engine ID is not configured first, the configuration command will fail. For the *privpassword* and *auth password* arguments, the minimum length is one character; the recommended length is at least eight characters, and should include both letters and numbers. Cisco IOS SNMP Support Command Reference (2013), at 396	To configure a remote user, specify the IP address or port number of the device where the user's remote SNMP agent resides. A remote agent's engine ID must be configured before remote users for that agent are configured. A user's authentication and privacy digests are derived from the engine ID and the user's password. The configuration command fails if the remote engine ID is not configured first. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1999

Cisco IOS 15.4
 	
Effective date of registration: 11/26/2014
	timers basic (ISO CLNS)
	
	To configure ISO IGRP timers, use the **timers basic** command in router configuration mode. To restore the default values, use the **no** form of this command.
	
	timers basic *update interval holddown interval invalid interval*
	
	no timers basic *update interval holddown-interval invalid-interval*
	
	Syntax Description
	update-interval — Time, in seconds, between the sending of routing updates.
	holddown-interval — Time, in seconds, a system or area router is kept in holddown state, during which routing information regarding better paths is suppressed. (A router enters into a holddown state when an update packet is received that indicates the route is unreachable. The route is marked inaccessible and advertised as unreachable. However, the route is still used for forwarding packets.) When the holddown interval expires, routes advertised by other sources are accepted and the route is no longer inaccessible.
	invalid-interval — Time, in seconds, that a route remains in the routing table after it has been determined that it is not reachable. After that length of time, the route is removed from the routing table.
	
	Cisco IOS Interface and Hardware Component Command Reference (2011), at ISO-178.
		timers basic (RIP)
		
		The timers basic command configures the update interval, the expiration time, and the deletion time for routes received and sent through RIP. The command requires value declaration of all values.
		
		• The update time is the interval between unsolicited route responses. The default is 30 seconds.
		• The expiration time is initialized when a route is established and any time an update is received for the route. If the specified period elapses from the last time the route update was received, then the route is marked as inaccessible and advertised as unreachable. However, the route forwards packets until the deletion time expires. The default value is 180 seconds.
		• The deletion time is initialized when the expiration time has elapsed. On initialization of the deletion time, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped. Upon expiration of the deletion time, the route is removed from the routing table. The default is 120 seconds.
		
		The no timers basic and default timers basic commands return the timer values to their default values by removing the timers-basic command from *running-config*.
		
		Platform — all
		Command Mode — Router-RIP Configuration
		
		Command Syntax
		`timers basic update time expire_time deletion_time`
		`no timers basic`
		`default timers basic`
		
		Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1671

Copyright Registration Information	Cisco	Arista
Cisco IOS 15.4 Effective date of registration: 11/26/2014	 **Field** **Description** Version 34 Indicates version number of the Level 1 routing table. All Level 1 routes with a version number that does not match this number are flushed from the routing table. The router's version number increments when the configuration changes from Level 1 or Level 1-2 to Level 2 only. System Id Identification value of the system listed in Level 1 forwarding table. Next-Hop System ID of best-cost next-hop to listed address. SNPA SNPA of next-hop system. Interface Interface through which next-hop system is known. Metric IS-IS metric for the route. State Up (active) or Down (nonoperational). Cisco IOS Interface and Hardware Component Command Reference (2011), at ISO-137.	**Display Values** • Inst. ID IS-IS Instance name. • System ID Identification value of the system listed in the Level 2 forwarding table. • Type Level 2 information. • Interface Interface through which the neighbor is reachable. • SNPA Subnetwork point of attachment (MAC address of the next hop). • State State of the adjacency: Up, Down, or INIT • Hold time Remaining hold time of the adjacency. • Area Address The address of the area. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1702
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Building the Address Table and Address Table Changes** The device dynamically builds the address table by using the MAC source address of the frames received. When the device receives a frame for a MAC destination address not listed in its address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the device adds its relevant MAC source address and port ID to the address table. The device then forwards subsequent frames to a single LAN port without flooding all LAN ports. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 10	**14.3 MAC Address Table** The switch maintains an MAC address table for switching frames efficiently between ports. The MAC address table contains static and dynamic MAC addresses. • Static MAC addresses are entered into the table through a CLI command. • Dynamic MAC addresses are entered into the table when the switch receives a frame whose source address is not listed in the MAC address table. The switch builds the table dynamically by referencing the source address of frames it receives. When the switch receives a frame, it associates the MAC address of the transmitting interface with the recipient VLAN. When a VLAN receives a frame for a MAC destination address not listed in the address table, the switch bridges the frame to all of the VLAN's ports except the recipient port. When the destination interface replies, the switch adds its MAC address to the MAC address table. The switch forwards subsequent frames with the destination address to the specified port. A multicast address can be associated with multiple ports. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 624

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	• Community VLAN—A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port gateways and to other host ports in the same community. You can configure multiple community VLANs in a private VLAN domain. The ports within one community can communicate, but these ports cannot communicate with ports in any other community or isolated VLAN in the private VLAN. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 54	— Community Community VLAN ports carry traffic from host ports to the primary VLAN ports and to other host ports in the same community VLAN. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 763
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	• Protocol migration—For backward compatibility with 802.1D devices, 802.1w selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis. When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which 802.1w BPDUs are sent), and 802.1w BPDUs are sent. While this timer is active, the device processes all BPDUs received on that port and ignores the protocol type. If the device receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D device and starts using only 802.1D BPDUs. However, if the 802.1w device is using 802.1D BPDUs on a port and receives an 802.1w BPDU after the timer has expired, it restarts the timer and starts using 802.1w BPDUs on that port. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 100	The clear spanning-tree detected-protocols command forces MST ports to renegotiate with their neighbors. RSTP provides backward compatibility with 802.1D bridges as follows: • RSTP selectively sends 802.1D-configured BPDUs and Topology Change Notification (TCN) BPDUs on a per-port basis. • When a port initializes, the migration delay timer starts and RSTP BPDUs are transmitted. While the migration delay timer is active, the bridge processes all BPDUs received on that port. • If the bridge receives an 802.1D BPDU after a port's migration delay timer expires, the bridge assumes it is connected to an 802.1D bridge and starts using only 802.1D BPDUs. • When RSTP uses 802.1D BPDUs on a port and receives an RSTP BPDU after the migration delay expires, RSTP restarts the migration delay timer and resumes using RSTP BPDUs on that port. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 953
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Loop Guard** Loop Guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 176	• Loop Guard: Prevents loops resulting from a unidirectional link failure on a point-to-point link. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 963
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Rapid PVST+ achieves rapid transition to the forwarding state only on edge ports and point-to-point links. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 90	RSTP only achieves rapid transition to forwarding state on edge ports and point-to-point links. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 964

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Enabling Loop Guard on a root device has no effect but provides protection when a root device becomes a nonroot device. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 176	Enabling loop guard on a root switch has no effect until the switch becomes a nonroot switch. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 966
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	• Enabling Loop Guard globally works only on point-to-point links. • Enabling Loop Guard per interface works on both shared and point-to-point links. • Root Guard forces a port to always be a designated port; it does not allow a port to become a root port. Loop Guard is effective only if the port is a root port or an alternate port. You cannot enable Loop Guard and Root Guard on a port at the same time. • Loop Guard has no effect on a disabled spanning tree instance or a VLAN. • Spanning tree always chooses the first operational port in the channel to send the BPDUs. If that link becomes unidirectional, Loop Guard blocks the channel, even if other links in the channel are functioning properly. • If you group a set of ports that are already blocked by Loop Guard to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role. • If a channel is blocked by Loop Guard and the channel members go back to an individual link status, spanning tree loses all the state information. The individual physical ports may obtain the forwarding state with the designated role, even if one or more of the links that formed the channel are unidirectional. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 179	Loop guard, when enabled globally, applies to all point-to-point ports. Loop guard is configurable on individual ports and applies to all STP instances of an enabled port. Loop-inconsistent ports transition to listening state when loop guard is disabled. Enabling loop guard on a root switch has no effect until the switch becomes a nonroot switch. When using loop guard: • Do not enable loop guard on portfast-enabled ports. • Loop guard is not functional on ports not connected to point-to-point links. • Loop guard has no effect on disabled spanning tree instances. Loop guard aspects on port channels include: • BPDUs are sent over the channel's first operational port. Loop guard blocks the channel if that link becomes unidirectional even when other channel links function properly. • Creating a new channel destroys state information for its component ports; new channels with loop-guard-enabled ports can enter forwarding state as a DP. • Dissembling a channel destroys its state information; component ports from a blocked channel can enter the forwarding state as DPs, even if the channel contained unidirectional links. • A unidirectional link on any port of a loop-guard-enabled channel blocks the entire channel until the affected port is removed or the link resumes bidirectional operation. Loop guard configuration commands include: • spanning-tree loopguard default command enables loop guard as a default on all switch ports. • spanning-tree guard control the loop guard setting on the configuration mode interface. This command overrides the default command for the specified interface. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 966

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**BPDU Guard** Enabling BPDU Guard shuts down that interface if a BPDU is received You can configure BPDU Guard at the interface level. When configured at the interface level, BPDU Guard shuts the port down as soon as the port receives a BPDU, regardless of the port type configuration. When you configure BPDU Guard globally, it is effective only on operational spanning tree edge ports. In a valid configuration, Layer 2 LAN edge interfaces do not receive BPDUs. A BPDU that is received by an edge Layer 2 LAN interface signals an invalid configuration, such as the connection of an unauthorized device. BPDU Guard, when enabled globally, shuts down all spanning tree edge ports when they receive a BPDU BPDU Guard provides a secure response to invalid configurations, because you must manually put the Layer 2 LAN interface back in service after an invalid configuration Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 174-75	20.3.4.3 BPDU Guard PortFast interfaces do not receive BPDUs in a valid configuration. BPDU Guard provides a secure response to invalid configurations by disabling ports when they receive a BPDU. Disabled ports differ from blocked ports in that they are re-enabled only through manual intervention. • When configured globally, BPDU Guard is enabled on ports in the operational portfast state. • When configured on an individual interface, BPDU Guard disables the port when it receives a BPDU, regardless of the port's portfast state. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 968
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**BPDU Filtering** You can use BPDU Filtering to prevent the device from sending or even receiving BPDUs on specified ports. When configured globally, BPDU Filtering applies to all operational spanning tree edge ports. You should connect edge ports only to hosts, which typically drop BPDUs. If an operational spanning tree edge port receives a BPDU, it immediately returns to a normal spanning tree port type and moves through the regular transitions. In that case, BPDU Filtering is disabled on this port, and spanning tree resumes sending BPDUs on this port In addition, you can configure BPDU Filtering by the individual interface. When you explicitly configure BPDU Filtering on a port, that port does not send any BPDUs and drops all BPDUs that it receives. You can effectively override the global BPDU Filtering setting on individual ports by configuring the specific interface. This BPDU Filtering command on the interface applies to the entire interface, whether the interface is trunking or not Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 175	20.3.4.4 BPDU Filter BPDU filtering prevents the switch from sending or receiving BPDUs on specified ports. BPDU filtering is configurable on Ethernet and port channel interfaces. Ports with BPDU filtering enabled do not send BPDUs and drops inbound BPDUs. Enabling BPDU filtering on a port not connected to a host can result in loops as the port continues forwarding data while ignoring inbound BPDU packets. The spanning-tree bpdufilter command controls BPDU filtering on the configuration mode interface. BPDU filtering is disabled by default. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 968

CASE NO. 5:14-CV-05344-BLF
EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Bridge Assurance** You can use Bridge Assurance to protect against certain problems that can cause bridging loops in the network. Specifically, you use Bridge Assurance to protect against a unidirectional link failure or other software failure and a device that continues to forward data traffic when it is no longer running the spanning tree algorithm. Note: Bridge Assurance is supported only by Rapid PVST+ and MST. Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 175	**spanning-tree bridge assurance** The spanning-tree bridge assurance command enables bridge assurance on all ports with a port type of *network*. Bridge assurance protects against unidirectional link failure, other software failure, and devices that quit running a spanning tree algorithm. Bridge assurance is available only on spanning tree *network* ports on point-to-point links. Both ends of the link must have bridge assurance enabled. If the device on one side of the link has bridge assurance enabled and the device on the other side either does not support bridge assurance or does not have it enabled, the bridge assurance enabled port is blocked. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1002
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	• Root Guard—Root Guard prevents the port from becoming the root in an STP topology. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 6	• Root guard prevents a port from becoming a root or blocked port. A root guard port that receives a superior BPDU transitions to the root-inconsistent (blocked) state. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1005
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Note: Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the VLAN. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 108	Important: When disabling spanning tree on a VLAN, ensure that all switches and bridges in the network disable spanning tree for the same VLAN. Disabling spanning tree on a subset of switches and bridges in a VLAN may have unexpected results because switches and bridges running spanning tree will have incomplete information regarding the network's physical topology. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1023
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	The software elects a router as the IGMP querier on a subnet if it has the lowest IP address. As long as a router continues to receive query messages from a router with a lower IP address, it resets a timer that is based on its querier timeout value. If the querier timer of a router expires, it becomes the designated querier. If that router later receives a host query message from a router with a lower IP address, it drops its role as the designated querier and sets its querier timer again. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 20	The router with the lowest IP address on a subnet sends membership queries as the IGMP querier. When a router receives a membership query from a source with a lower IP address, it resets its query response timer. Upon timer expiry, the router begins sending membership queries. If the router subsequently receives a membership query from a router with a lower IP address, it stops sending membership queries and resets the query response timer. Arista User Manual v. 4v. 4.14.3F - Rev. 2 (10/2/14), at 1779

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	IGMP version: 2 Startup query interval: 30 seconds Startup query count: 2 Robustness value: 2 Querier timeout: 255 seconds Query timeout: 255 seconds Query max response time: 10 seconds Query interval: 125 seconds Last member query response interval: 1 second Last member query count: 2 Group membership timeout: 260 seconds Report link local multicast groups: Disabled Enforce router alert: Disabled Immediate leave: Disabled Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 24	Current IGMP router version: 2 IGMP query interval: 125 seconds IGMP max query response time: 100 deciseconds Last member query response interval: 10 deciseconds Last member query response count: 2 IGMP querier: 172.17.26.1 Robustness: 2 Require router alert: enabled Startup query interval: 312 deciseconds Startup query count: 2 General query timer expiry: 00:00:22 Multicast groups joined: 239.255.255.250 Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1850

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Anycast-RP** Anycast-RP has two implementations: one uses Multicast Source Discovery Protocol (MSDP) and the other is based on *RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM)*. This section describes how to configure PIM Anycast-RP. You can use PIM Anycast-RP to assign a group of routers, called the Anycast-RP set, to a single RP address that is configured on multiple routers. The set of routers that you configure as Anycast-RPs is called the Anycast-RP set. This method is the only RP method that supports more than one RP per multicast group, which allows you to load balance across all RPs in the set. The Anycast RP supports all multicast groups. PIM register messages are sent to the closest RP and PIM join-prune messages are sent in the direction of the closest RP as determined by the unicast routing protocols. If one of the RPs goes down, unicast routing ensures these message will be sent in the direction of the next-closest RP. You must configue PIM on the loopback interface that is used for the PIM Anycast RP. For more information about PIM Anycast-RP, see *RFC 4610*. For information about configuring Anycast-RPs, see *Configuring a PIM Anycast-RP Set*. **PIM Register Messages** PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The PIM register message has the following functions: • To notify the RP that a source is actively sending to a multicast group. • To deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues to send PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 68-69	**Anycast-RP** PIM Anycast-RP defines a single RP address that is configured on multiple routers. An anycast-RP set consists of the routers configured with the same anycast-RP address. Anycast-RP provides redundancy protection and load balancing. The anycast-RP set supports all multicast groups. PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The switch sends these messages and join-prune messages to the anycast-RP set member specified in the anycast-RP command. In a typical configuration, one command is required for each member of the anycast-RP set. The PIM register message has the following functions: • Notify the RP that a source is actively sending to a multicast group. • Deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues sending PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1874
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Note** Use the show ip mroute command to display the statistics for multicast route and prefixes. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 118	**Multicast Display Commands** To display the information in the multicast routing table use the show ip mroute command. To display the MFIB table information, use the show ip mfib command. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1758

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	 Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 139	The ip igmp snooping command controls the global snooping setting. The ip igmp snooping vlan command enables snooping on individual VLANs if snooping is globally enabled. IGMP snooping is enabled on all VLANs by default. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1780
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	 Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 140	**Specifying a Static Multicast Router Connection** The ip igmp snooping vlan mrouter command statically configures a port that connects to a multicast router to join all multicast groups. The port to the router must be in the specified VLAN range. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1780
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Displaying IGMP Snooping Statistics** Use the show ip igmp snooping statistics vlan command to display IGMP snooping statistics. You can see the virtual port channel (vPC) statistics in this output. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 144	**show ip igmp statistics** The show ip igmp statistics command displays IGMP transmission statistics for the specified interface. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1867

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**SA Messages and Caching** MSDP peers exchange Source-Active (SA) messages to propagate information about active sources. SA messages contain the following information: • Source address of the data source. • Group address that the data source uses • IP address of the RP or the configured originator ID Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 148-49	35.2.2.1 Source Active Messages A Source Active (SA) message is a message that an RP creates and sends to MSDP peers when it learns of a new multicast source through a PIM register message. RPs that intend to originate or receive SA messages must establish MSDP peering with other RPs, either directly or through intermediate MSDP peers. An RP that is not a DR on a shared network should only originate SAs in response to register messages it receives from the DR. It does not originate SA's for directly connected sources in its domain. SA messages contain the following fields: • Source address of the data source. • Group address that receives data sent by the source. • IP address of the RP. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1912
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	RFC 5059 *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)* Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 174	**34.3 Configuring PIM** The following sections describe the configuration of static RPs, dynamic RPs, and anycast-RPs. RP implementation is defined through the following RFCs: • RFC 5059: Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM). • RFC 6226: PIM Group-to-Rendezvous-Point Mapping. This section describes the following configuration tasks: • Section 34.3.1: Enabling PIM • Section 34.3.2: Rendezvous Points (RPs) • Section 34.3.3: Hello Messages • Section 34.3.4: Designated Router Election • Section 34.3.5: Join-Prune Messages Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1872
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Audience** This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches. Cisco DCNM Fundamentals Guide, Release 6.x (2011), at lxi	**Audience** This guide is for experienced network administrators who are responsible for configuring and maintaining Arista switches. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 41

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Table 5-1 Channel Modes for Individual Links in a Port Channel** | Channel Mode	Description
---	---
passive	LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.
active	LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets.
on	All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the on state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. | Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 5-10	**Parameters** • *number* specifies a channel group ID. Values range from 1 through 1000. • *LACP_MODE* specifies the interface LACP mode. Values include: — mode **on** Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode **active** Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode **passive** Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 469
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Note** For information about configuring port channels and the Link Aggregation Control Protocol (LACP), see Chapter 5, "Configuring Port Channels." Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 6-2	# Port Channels and LACP This chapter describes channel groups, port channels, port channel interfaces, and the Link Aggregation Control Protocol (LACP). This chapter contains the following sections: Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 469
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Loopback Interfaces** A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 4-4 .	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 631

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Configuring a Maximum Number of MAC Addresses** You can configure the maximum number of MAC addresses that can be learned or statically configured on interfaces that belong to a port profile. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 10-22	**Port Security Configuration** MAC security restricts input to a switched port by limiting the number and identity of MAC addresses that can access the port. MAC address security is enabled by switchport port-security. Ports with MAC security enabled restrict traffic to a limited number of hosts, as determined by their MAC addresses. The maximum number of MAC addresses that can be assigned to an interface is configured by switchport port-security maximum. The default MAC address limit on an interface where port security is enabled is one. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 632

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to return to EXEC mode from global configuration mode:
	
	`switch(config)# end`
	`switch#`
	
	This example shows how to return to EXEC mode from interface configuration mode:
	
	`switch(config-if)# end`
	`switch#`
	
	Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-44
		• To return to Privileged EXEC mode from any configuration mode, type end or Ctrl-Z.
		
		`switch(config-if-Et24)#<Ctrl-z>`
		`switch#`
		
		Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 120

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Note The reload command does not save the running configuration. Use the copy running-config startup-config command to save the current configuration on the device.
	
	Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-105
		Step 8 Type write memory (or copy running-config startup-config) to save the new configuration to the *startup-config* file. See Section 3.5.4: Saving the Running Configuration Settings.
		
		`switch# write memory`
		`switch#`
		
		Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 60

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to display commands related to Open Shortest Path First (OSPF) available in the loopback interface command mode:
	
	`switch(config)# interface loopback 0`
	`switch(config-if)# show cli list ospf`
	`MODE if-loopback`
	`no ip ospf network point-to-point`
	`no ip ospf network`
	
	Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-126
		Command Syntax
		`ip ospf network point-to-point`
		`no ip ospf network`
		`default ip ospf network`
		
		Examples
		• These commands configure Ethernet interface 10 as a point-to-point link.
		
		`switch(config)#interface ethernet 10`
		`switch(config-if-Et10)#ip ospf network point-to-point`
		`switch(config-if-Et10)#`
		
		• This command restores Ethernet interface 10 as a broadcast link.
		
		`switch(config-if-Et10)#no ip ospf network`
		`switch(config-if-Et10)#`
		
		Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 1432

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	show startup-config
	
	To display the startup configuration, use the show **startup-config** command.
	
	show startup-config [exclude *component-list*]
	
	Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-154
		Example
		• Type show startup-config to display the startup configuration file. The response in the example is truncated to display only the ip route configured in Admin Username (page 58).
		
		`switch#show startup-config`
		`! Command: show startup-config`
		`! Startup-config last modified at  Wed Feb 19 08:34:31 2014 by admin`
		`!`
		`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
		`!`
		`ip route 0.0.0.0/0 192.0.2.1`
		`!`
		`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
		`end`
		`switch#`
		
		Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 123

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Enabling the Error-Disable Detection** You can enable error-disable detection in an application. As a result, when a cause is detected on an interface, the interface is placed in an error disabled state, which is an operational state that is similar to the link-down state. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 2-24	**14.5.2 Errdiabled Ports** The switch places an Ethernet or management interface in *error-disabled* state when it detects an error on the interface. *Error-disabled* is an operational state that is similar to link-down state. Conditions that error-disables an interface includes: Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 123

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	This example shows how to configure a Layer 2 trunk interface, assign the native VLAN and the allowed VLANs, and configure the device to tag the native VLAN traffic on the trunk interface: switch# configure terminal switch(config)# interface ethernet 2/35 switch(config-if)# switchport switch(config-if)# switchport mode trunk switch(config-if)# switchport trunk native vlan 10 switch(config-if)# switchport trunk allowed vlan 5, 10 switch(config-if)# exit switch(config)# vlan dot1q tag native switch(config)# Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 3-36	The trunk group command is not additive to the allowed vlan command interface ethernet 1 switchport mode trunk switchport trunk allowed vlan 10 switchport trunk group trunk30 Vlan 30 will not be permitted on the interface as it is not listed in the allowed vlan list. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 767
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Loopback Interfaces** A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 4-4	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 631
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	end **Example:** switch(config-router-af)# end Exits address family configuration mode and returns to global configuration mode. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 5-30	• This command exits server-failure configuration mode and returns to global configuration mode. switch(config-server-failure)#exit switch(config)# Arista User Manual v. 4.14.3F - Rev. 2 (10/2/14), at 640

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Table 6-1 Channel Modes for Individual Links in a Port Channel** | Channel Mode	Description
---	---
passive	LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.
active	LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets.
on	All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. | Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 6-10	• LACP_MODE specifies the interface LACP mode. Values include: — mode on Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode active Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode passive Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 469
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Configuring the LACP Fast Timer Rate** You can change the LACP timer rate to modify the duration of the LACP timeout. Use the lacp rate command to set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to the fast rate (1 second). This command is supported only on LACP-enabled interfaces. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 6-38	**lacp rate** The lacp rate command configures the LACP transmission interval on the configuration mode interface. The LACP timeout sets the rate at which LACP control packets are sent to an LACP-supported interface. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 478
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Step 3 `lacp rate fast` Configures the fast rate (one second) at which LACP control packets are sent to an LACP-supported interface. Example: `switch(config-if)# lacp rate fast` To reset the timeout rate to its default, use the **no** form of the command. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 6-38	**lacp rate** The lacp rate command configures the LACP transmission interval on the configuration mode interface. The LACP timeout sets the rate at which LACP control packets are sent to an LACP-supported interface. Supported values include: • *normal*: 30 seconds with synchronized interfaces; one second while interfaces are synchronizing. • *fast*: one second. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 478

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Syntax Description** ipv4 — (Optional) Configures BFD session parameters for the IPv4 address. ipv6 — (Optional) Configures BFD session parameters for the IPv6 address. *mintx* — Rate at which BFD control packets are sent to BFD neighbors. The configurable range is from 50 to 999. min_rx *msec* — Specifies the rate at which BFD control packets are expected to be received from BFD neighbors. The range is from 50 to 999. multiplier *value* — Specifies the number of consecutive BFD control packets that must be missed from a BFD neighbor before BFD declares that the neighbor is unavailable and the BFD neighbor is informed of the failure. The range is from 1 to 50. **Defaults** BFD interval: 50 milliseconds min_rx: 50 milliseconds multiplier: 3 Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 1-12	31.3.1 Configuring BFD on an Interface The transmission rate for BFD control packets, the minimum rate at which control packets are expected from the peer, and the multiplier (the number of packets that must be missed in succession before BFD declares the session to be down) are all configured per interface. These values apply to all BFD sessions that pass through the interface. The default values for these parameters are: • transmission rate 300 milliseconds • minimum receive rate 300 milliseconds • multiplier 3 Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 1737
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**ip pim bfd-instance** To enable Bidirectional Forwarding Detection (BFD) for Protocol Independent Multicast (PIM) on an interface, use the **ip pim bfd-instance** command. To return to the default setting, use the no form of this command. ip pim bfd-instance [disable] no ip pim bfd-instance [disable] Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 1-251	31.3.2 Configuring BFD for PIM To enable or disable bidirectional forwarding detection (BFD) globally for all protocol independent multicast (PIM) neighbors, use the ip pim bfd command. To enable or disable PIM BFD on a specific interface, use the ip pim bfd-instance command. The interface-level configuration supercedes the global setting. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 766
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**switchport trunk allowed vlan** To set the list of allowed VLANs on the trunking interface, use the switchport trunk allowed vlan command. To allow all VLANs on the trunking interface, use the no form of this command. switchport trunk allowed vlan {vlan-list	add vlan-list	all	except vlan-list	none	remove vlan-list} no switchport trunk allowed vlan Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 1-251	By default all VLAN's are permitted on a port configured with 'switchport mode trunk'. To limit the port's VLAN trunk list use the switchport trunk allowed vlan command. Only VLAN's in the allowed list will be permitted. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 766

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	switchport trunk native vlan
	
	To change the native VLAN ID when the interface is in trunking mode, use the **switchport trunk native vlan** command. To return the native VLAN ID to VLAN 1, use the **no form** of this command.
	
	switchport trunk native vlan *vlan-id*
	
	no switchport trunk native vlan
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 1-253
		To specify the port's native VLAN, use the switchport trunk native vlan command.
		
		Example
		
		• These commands configure VLAN 12 as the native VLAN trunk for Ethernet interface 10.
		
		`switch(config)#interface ethernet 10`
		`switch(config-if-Et10)#switchport trunk native vlan 12`
		`switch(config-if-Et10)#`
		
		Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 766

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to clear all the dynamic Layer 2 entries from the MAC address table for VLAN 20 on port 2/20:
	
	`switch(config)# clear mac address-table dynamic vlan 20 interface ethernet 2/20`
	`switch(config)#`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 3
		Example
		
		• This command clears all dynamic mac address table entries for port channel 5 on VLAN 34.
		
		`switch#clear mac address-table dynamic vlan 34 interface port-channel 5`
		`switch#`
		
		Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 648

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Usage Guidelines** Rapid per VLAN Spanning Tree Plus (Rapid PVST+) and Multiple Spanning Tree (MST) have built-in compatibility mechanisms that allow them to interact properly with other versions of IEEE spanning tree or other regions. For example, a bridge running Rapid PVST+ can send 802.1D bridge protocol data units (BPDUs) on one of its ports when it is connected to a legacy bridge. An MST bridge can detect that a port is at the boundary of a region when it receives a legacy BPDU or an MST BPDU that is associated with a different region. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 5	20.2.1.4 Version Interoperability A network can contain switches running different spanning tree versions. The common spanning tree (CST) is a single forwarding path the switch calculates for STP, RSTP, MSTP, and Rapid-PVST topologies in networks containing multiple spanning tree variations. In multi-instance topologies, the following instances correspond to the CST: • Rapid-PVST: VLAN 1 • MST: IST (instance 0) RSTP and MSTP are compatible with other spanning tree versions: • An RSTP bridge sends 802.1D (original STP) BPDUs on ports connected to an STP bridge. • RSTP bridges operating in 802.1D mode remain in 802.1D mode even after all STP bridges are removed from their links. • An MST bridge can detect that a port is at a region boundary when it receives an STP BPDU or an MST BPDU from a different region. • MST ports assume they are boundary ports when the bridges to which they connect join the same region. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 953

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Examples This example shows how to add a static entry to the MAC address table.
	`switch(config)+ mac address-table static 0050.3e9d.6400 vlan 3 interface ethernet 2/1`
	`switch(config)#`
	
	Related Commands
		Command	Description	
				show mac address-table
					Displays information about the MAC address table.
					
					Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 20
						The mac address-table static command adds a static entry to the MAC address table.
						
						Example
						• This command adds a static entry for unicast MAC address 0012.3694.03ec to the MAC address table.
						`switch(config)#mac address-table static 0012.3694.03ec vlan 3 interface Ethernet 7`
						`switch(config)#show mac address-table static`
						
						Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 624

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Related Commands**	Command	Description | show spanning-tree mst configuration	Displays information about the MST protocol. spanning-tree mst configuration	Enters MST configuration submode. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 24	**show spanning-tree mst configuration** The show spanning-tree mst configuration command displays information about the MST region's VLAN-to-instance mapping. The command provides two display options: • default displays a table that lists the instance to VLAN map. • digest displays the configuration digest. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 991

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Examples This example shows how to display VTP interface switchport information on the device:
	`switch# show interface switchport`
	`Name: Ethernet8/11`
	`  Switchport: Enabled`
	`  Switchport Monitor: Not enabled`
	`  Operational Mode: trunk`
	`  Access Mode VLAN: 1 (default)`
	`  Trunking Native Mode VLAN: 1 (default)`
	`  Trunking VLANs Enabled: 1,10,20-30`
	`  Pruning VLANs Enabled: 2-1001`
	`  Administrative private-vlan primary host-association: none`
	`  Administrative private-vlan secondary host-association: none`
	`  Administrative private-vlan primary mapping: none`
	`  Administrative private-vlan secondary mapping: none`
	`  Administrative private-vlan trunk native VLAN: none`
	`  Administrative private-vlan trunk encapsulation: dot1q`
	`  Administrative private-vlan trunk normal VLANs: none`
	`  Administrative private-vlan trunk private VLANs: none`
	`  Operational private-vlan: none`
	`switch#`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 44
		Example
		• These commands create the trunk mode allowed VLAN list of 6-10 for Ethernet interface 14, then verifies the VLAN list.
		`switch(config)#interface ethernet 14`
		`switch(config-if-Et14)#switchport trunk allowed vlan 6-10`
		`switch(config-if-Et14)#show interfaces ethernet 14 switchport`
		`Name: Et14`
		`  Switchport: Enabled`
		`  Administrative Mode: trunk`
		`  Operational Mode: trunk`
		`  Access Mode VLAN: 1 (inactive)`
		`  Trunking Native Mode VLAN: 1 (inactive)`
		`  Administrative Native VLAN tagging: disabled`
		`  Trunking VLANs Enabled: 6-10`
		`  Trunk Groups:`
		
		`switch(config-if-Et14)#`
		
		Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 798

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Examples This example shows how to display information about the specified VLAN. This command displays statistical information gathered on the VLAN at 1-minute intervals:
	
	`switch# show interface vlan 5`
	`Vlan5 is administratively down, line protocol is down`
	`  Hardware is EtherSVI, address is  0000.0000.0000`
	`  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,`
	`    reliability 255/255, txload 1/255, rxload 1/255`
	`  Encapsulation ARPA, loopback not set`
	`  Keepalive not supported`
	`  ARP type: ARPA`
	`  Last clearing of "show interface" counters 01:21:55`
	`  1 minute input rate 0 bytes/sec, 0 packets/sec`
	`  1 minute output rate 0 bytes/sec, 0 packets/sec`
	`  L3 Switched:`
	`    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes`
	`  L3 in Switched:`
	`    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes`
	`  L3 out Switched:`
	`    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 44
		Example
		• This command display configuration and status information for Ethernet interface 1 and 2.
		
		`switch>show interfaces ethernet 1-2`
		`Ethernet1 is up, line protocol is up (connected)`
		`  Hardware is Ethernet, address is 001c.2481.7647 (bia 001c.2481.7647)`
		`  Description: mkt.1`
		`  MTU 9212 bytes, BW 10000000 Kbit`
		`  Full-duplex, 10Gb/s, auto negotiation: off`
		`  Last clearing of "show interface" counters never`
		`  5 seconds input rate 33.5 Mbps (0.3% with framing), 846 packets/sec`
		`  5 seconds output rate 180 kbps (0.0% with framing), 55 packets/sec`
		`     76437268 packets input, 94280286608 bytes`
		`     Received 2208 broadcasts, 73358 multicast`
		`     0 runts, 0 giants`
		`     0 input errors, 0 CRC, 0 alignment, 0 symbol`
		`     0 PAUSE input`
		`     6184281 packets output, 4071319140 bytes`
		`     Sent 2209 broadcasts, 345754 multicast`
		`     0 output errors, 0 collisions`
		`     0 late collision, 0 deferred`
		`     0 PAUSE output`
		
		Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 437

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**show mac address-table** To display the information about the MAC address table, use the show mac address-table command. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 54	**14.3.2 Displaying the MAC Address Table** The show mac address-table command displays the specified MAC address table entries. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 626
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Command** / **Description** mac address-table static — Adds static entries to the MAC address table or configures a static MAC address with IGMP snooping disabled for that address. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 56	**mac address-table static** The mac address-table static command adds a static entry to the MAC address table. Each table entry references a MAC address, a VLAN, and a list of layer 2 (Ethernet or port channel) ports. The table supports three entry types: unicast drop, unicast, and multicast. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 664

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Command** **Description** mac address-table aging-time — Configures the aging time for entries in the Layer 2 table. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 57	The mac address-table aging-time command configures the aging time for MAC address table dynamic entries. Aging time defines the period an entry is in the table, as measured from the most recent reception of a frame on the entry's VLAN from the specified MAC address. The switch removes entries when their presence in the MAC address table exceeds the aging time. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 626
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Examples** This example shows how to display STP when you are running Rapid PVST+: ``` switch# show spanning-tree VLAN0001 Spanning tree enabled protocol rstp Root ID Priority 32769 Address 000d.eca3.9f01 Cost 4 Port 4105 (port-channel10) Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 32769 (priority 32768 sys-id-ext 1) Address 0022.5579.7641 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Interface Role Sts Cost Prio.Nbr Type Po10 Root FWD 2 128.4105 (vPC peer-link) P2p Po20 Desg FWD 1 128.4115 (vPC) P2p Po30 Root FWD 1 128.4125 (vPC) P2p ``` Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 63	Show commands (such as show spanning-tree) displays the RSTP instance as MST0 (MST instance 0). **Example** • This command, while the switch is in RST mode, displays RST instance information. ``` switch(config)#show spanning-tree MST0 Spanning tree enabled protocol rstp <---RSTP mode indicator Root ID Priority 32768 Address 001c.730c.1867 This bridge is the root Bridge ID Priority 32768 (priority 32768 sys-id-ext 0) Address 001c.730c.1867 Hello Time 2.000 sec Max Age 20 sec Forward Delay 15 sec Interface Role State Cost Prio.Nbr Type --------------- --------- --------- -------- -------------- Et51 designated forwarding 2000 128.51 P2p switch(config)# ``` Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 960

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	This example shows how to display STP information when you are running MST: ``` switch# show spanning-tree MST0000 Spanning tree enabled protocol mstp Root ID Priority 32768 Address 0018.bad8.fc150 Cost 0 Port 258 (Ethernet 2/2) Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 32768 (priority 32768 sys-id-ext 0) Address 0018.bad8.239d Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Interface Role Sts Cost Prio.Nbr Type ---------------- ---- --- --------- -------- -------------------------- Eth2/1 Altn BKN 20000 128.257 Network, P2p BA_Inc. Eth2/2 Root FWD 20000 128.258 Edge, P2p Eth3/48 Desg FWD 20000 128.43228 P2p ``` Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 63	This command displays output from the show spanning-tree command: ``` Switch#show spanning-tree MST0 Spanning tree enabled protocol mstp Root ID Priority 32768 Address 0011.2201.0301 This bridge is the root Bridge ID Priority 32768 (priority 32768 sys-id-ext 0) Address 0011.2201.0301 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Interface Role State Cost Prio.Nbr Type ---------------- ----------- ----------- --------- -------- ----------- Et4 designated forwarding 2000 128.4 P2p Et5 designated forwarding 2000 128.5 P2p ... PEt4 designated forwarding 2000 128.31 P2p PEt5 designated forwarding 2000 128.44 P2p ... Po3 designated forwarding 1999 128.1003 P2p ``` Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 983
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	``` Spanning tree enabled protocol rstp Root ID Priority 32770 Address 000d.eca3.9f01 Cost 4 Port 4105 (port-channel10) Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 32770 (priority 32768 sys-id-ext 2) Address 0022.5579.7641 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Interface Role Sts Cost Prio.Nbr Type -------------- ---- --- --------- -------- -------------------------- Po10 Root FWD 2 128.4105 (vPC peer-link) P2p Po20 Desg FWD 1 128.4115 (vPC) P2p Po30 Root FWD 1 128.4125 (vPC) P2p ``` Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 73	``` Spanning tree enabled protocol rstp Root ID Priority 32768 Address 001c.7301.07b9 Cost 1999 (Ext) 0 (Int) Port 101 (Port-Channel2) Hello Time 2.000 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 32768 (priority 32768 sys-id-ext 0) Address 001c.7304.195b Hello Time 2.000 sec Max Age 20 sec Forward Delay 15 sec Interface Role State Cost Prio.Nbr Type -------------- ----------- ----------- --------- -------- ----------- Et4 designated forwarding 20000 128.4 P2p Et5 designated forwarding 20000 128.5 P2p Et6 designated forwarding 20000 128.6 P2p Et23 designated forwarding 20000 128.23 P2p Et26 designated forwarding 20000 128.26 P2p Et32 designated forwarding 2000 128.32 P2p ``` Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 983

Copyright Registration Information	Cisco	Arista
	This example shows how to display detailed information about the STP configuration: ``` switch(config)# show spanning-tree detail VLAN0001 is executing the rstp compatible Spanning Tree protocol Bridge Identifier has priority 32768, sysid 1, address 0022.5579.7641 Configured hello time 2, max age 20, forward delay 15 Current root has priority 32769, address 000d.eca3.9f01 Root port is 4105 (port-channel10), cost of root path is 4 Topology change flag not set, detected flag not set Number of topology changes 1 last change occurred 20:24:36 ago from port-channel10 Times: hold 1, topology change 35, notification 2 hello 2, max age 20, forward delay 15 Timers: hello 0, topology change 0, notification 0 Port 4105 (port-channel10, vPC Peer-link) of VLAN0001 is root forwarding Port path cost 2, Port priority 128, Port Identifier 128.4105 Designated root has priority 32769, address 000d.eca3.9f01 Designated bridge has priority 32769, address 0022.5579.7341 Designated port id is 128.4105, designated path cost 2 Timers: message age 16, forward delay 0, hold 0 Number of transitions to forwarding state: 1 Link type is point-to-point by default BPDU: sent 36729, received 36739 Port 4115 (port-channel20, vPC) of VLAN0001 is designated forwarding Port path cost 1, Port priority 128, Port Identifier 128.4115 Designated root has priority 32769, address 000d.eca3.9f01 Designated bridge has priority 32769, address 0022.5579.7341 Designated port id is 128.4115, designated path cost 2 Timers: message age 0, forward delay 0, hold 0 Number of transitions to forwarding state: 0 Link type is point-to-point by default BPDU: sent 0, received 0 Port 4125 (port-channel30, vPC) of VLAN0001 is root forwarding Port path cost 1, Port priority 128, Port Identifier 128.4125 Designated root has priority 32769, address 000d.eca3.9f01 Designated bridge has priority 32769, address 000d.eca3.9f01 Designated port id is 128.4125, designated path cost 0 Timers: message age 0, forward delay 0, hold 0 Number of transitions to forwarding state: 0 Link type is point-to-point by default BPDU: sent 0, received 0 ``` Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 73	• This command displays STP data, including an information block for each interface running STP. ``` switch>show spanning-tree vlan 1000 detail MST0 is executing the rstp Spanning Tree protocol Bridge Identifier has priority 32768, sysid 0, address 001c.7304.195b Configured hello time 2.000, max age 20, forward delay 15, transmit hold-count 6 Current root has priority 32768, address 001c.7301.07b9 Root port is 101 (Port-Channel12), cost of root path is 1999 (Ext) 0 (Int) Number of topology changes 4109 last change occurred 1292651 seconds ago from Ethernet13 Port 4 (Ethernet4) of MST0 is designated forwarding Port path cost 20000, Port priority 128, Port Identifier 128.4. Designated root has priority 32768, address 001c.7301.07b9 Designated bridge has priority 32768, address 001c.7304.195b Designated port id is 128.4, designated path cost 1999 (Ext) 0 (Int) Timers: message age 1, forward delay 15, hold 20 Number of transitions to forwarding state: 1 Link type is point-to-point by default, Internal BPDU: sent 452252, received 0, taggedErr 0, otherErr 0, rateLimiterCount 0 Rate-Limiter: enabled, Window: 10 sec, Max-BPDU: 400 Port 5 (Ethernet5) of MST0 is designated forwarding Port path cost 20000, Port priority 128, Port Identifier 128.5. Designated root has priority 32768, address 001c.7301.07b9 Designated bridge has priority 32768, address 001c.7304.195b Designated port id is 128.5, designated path cost 1999 (Ext) 0 (Int) Timers: message age 1, forward delay 15, hold 20 Number of transitions to forwarding state: 1 Link type is point-to-point by default, Internal BPDU: sent 1006266, received 0, taggedErr 0, otherErr 0, rateLimiterCount 0 Rate-Limiter: enabled, Window: 10 sec, Max-BPDU: 400 ``` Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 984
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014		

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to display STP information about a specified interface when you are running Rapid PVST+:
	
	`switch(config)# show spanning-tree interface ethernet 8/2`
	
	`Vlan          Role Sts Cost      Prio.Nbr Type`
	`---------------- ---- --- --------- -------- --------------------------------`
	`VLAN0001      Altn BLK 20000    128.1025 P2p`
	
	`VLAN0002      Desg FWD 20000    128.1025 P2p`
	
	This example shows how to display STP information about a specified interface when you are running MST:
	
	`switch(config)# show spanning-tree interface ethernet 2/50`
	
	`Mst Instance    Role Sts Cost      Prio.Nbr Type`
	`---------------- ---- --- --------- -------- ---------------`
	`MST0000      Desg FWD 20000    128.1281 P2p`
	
	This example shows how to display detailed STP information about a specified interface when you are running Rapid PVST+:
	
	`switch(config)# show spanning-tree interface ethernet 8/1 detail`
	
	`Port 1025 (Ethernet8/1) of VLAN0001 is alternate blocking`
	`   Port path cost 20000, Port priority 128, Port Identifier 128.1025`
	`   Designated root has priority 28672, address 0018.bad8.239d`
	`   Designated bridge has priority 28672, address 0018.bad8.239d`
	`   Designated port id is 128.1281, designated path cost 0`
	`   Timers: message age 15, forward delay 0, hold 0`
	`   Number of transitions to forwarding state: 1`
	`   Link type is point-to-point by default`
	`   The port type is network by default.`
	`   BPDU: sent 4657, received 188`
	
	`Port 1025 (Ethernet8/1) of VLAN0002 is designated forwarding`
	`   Port path cost 20000, Port priority 128, Port Identifier 128.1025`
	`   Designated root has priority 32770, address 0018.bad7.fc15`
	`   Designated bridge has priority 32770, address 0018.bad7.fc15`
	`   Designated port id is 128.1025, designated path cost 0`
	`   Timers: message age 0, forward delay 0, hold 0`
	`   Number of transitions to forwarding state: 1`
	`   Link type is point-to-point by default`
	`   The port type is network by default.`
	`   BPDU: sent 4838, received 0`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 77.
		Examples
		• This command displays an STP table for Ethernet 5 interface.
		
		`switch>show spanning-tree interface ethernet 5`
		`Instance        Role      State     Cost      Prio.Nbr Type`
		`---------------- --------- --------- --------- -------- ---------------`
		`MST0            designated forwarding 20000    128.5    P2p`
		`switch>`
		
		• This command displays a data block for Ethernet interface 5.
		
		`switch>show spanning-tree interface ethernet 5 detail`
		`Port 5 (Ethernet5) of MST0 is designated forwarding`
		`   Port path cost 20000, Port priority 128, Port Identifier 128.5.`
		`   Designated root has priority 32768, address 001c.7301.07b9`
		`   Designated bridge has priority 32768, address 001c.7304.195b`
		`   Designated port id is 128.5, designated path cost 1999 (Ext) 0 (Int)`
		`   Timers: message age 1, forward delay 15, hold 20`
		`   Number of transitions to forwarding state: 1`
		`   Link type is point-to-point by default, Internal`
		`   BPDU: sent 1008766, received 0, taggedErr 0, otherErr 0, rateLimiterCount 0`
		`   Rate-Limiter: enabled, Window: 10 sec, Max-BPDU: 400`
		
		`switch>`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 988.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	 Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 80.	 Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 990.

Copyright Registration Information	Cisco	Arista
| Cisco NX-OS 6.2<br><br>Effective date of registration: 11/13/2014 | This example shows how to display information about the MST configuration:<br><br>`switch)# show spanning-tree mst configuration`<br><br>```
Name:        [mst-bldg-sj6/3]
Revision:   1         Instances Configured: 3
Instance    Vlans mapped
---------  ------------------------------------------
0          1
2000       2-2000
4094       2001-4094
------------------------------------------------------
```<br><br>This example shows how to display the MD5 digest included in the current MST configuration:<br><br>`switch)# show spanning-tree mst configuration digest`<br><br>```
Name      [mst-config]
Revision  10     Instances configured 25
Digest            0x40D5ECA178C657835C83BBCB16723192
Pre-std Digest   0x27BF112A75B72781ED928D9EC5BB4251
```<br><br>Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 81. | Examples<br>• This command displays the MST region's VLAN-to-instance map.<br><br>```
switch>show spanning-tree mst configuration
Name       []
Revision   0     Instances configured 3

Instance  Vlans mapped
--------  -------------------------------------------
0         1,4-4094
2         2
3         3
-----------------------------------------------------
switch>
```<br><br>• This command displays the MST region's configuration digest.<br><br>```
switch>show spanning-tree mst configuration digest
Name       []
Revision   0     Instances configured 1
Digest            0xAC36177F50283CD4B83821D8AB26DE62
switch>
```<br><br>Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 991. |
| Cisco NX-OS 6.2<br><br>Effective date of registration: 11/13/2014 | Examples     This example shows how to display information for the root bridge:<br><br>`switch(config)# show spanning-tree root`<br><br>```
MST Instance      Root ID         Cost  Time Age Dly  Root Port
-------------  ----------------  ----- ---- --- ---  ----------
MST0000         32768 0018.bad7.fc15   0    2   20  15  This bridge is root
```<br><br>Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 82-83. | Examples<br>• This command displays a table of root bridge information.<br><br>```
switch>show spanning-tree root
                        Root ID         Root  Hello Max Fwd
Instance     Priority   MAC addr    Cost  Time Age Dly  Root Port
----------  ----------  ----------- ----- ---- --- ---  ----------
MST0          32768 001c.7301.23de    0    2   20  15  Po937
MST101        32869 001c.7301.23de  3998   0    0   0  Po909
MST102        32870 001c.7301.23de  3998   0    0   0  Po911
switch>
```<br><br>Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 994. |

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to display information about the number of VLANs configured on the device:
	
	`switch# show vlan summary`
	
	`Number of existing VLANs      : 9`
	`  Number of existing user VLANs   : 9`
	`  Number of existing extended VLANs : 0`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 94.
		Example
		
		• This command displays the number of VLANs on the switch.
		
		`switch>show vlan summary`
		`Number of existing VLANs          : 18`
		
		`switch>`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 791.

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Examples This example shows how to display information about all private VLANs on the device:
	
	`switch(config)# show vlan private-vlan`
	
	`Primary  Secondary  Type          Ports`
	`-------  ---------  ----------    ------------------`
	`200      201        isolated      Eth2/26, Eth2/27`
	`200      202        community     Eth2/26, Eth2/28`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 100.
		Example
		
		• This command displays the private VLANs.
		
		`switch>show vlan private-vlan`
		`Primary  Secondary  Type          Ports`
		`-------  ---------  -----------   -----------------------------`
		`5        25         isolated`
		`5        26         isolated`
		`7        31         community`
		`7        32         isolated`
		`switch>`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 790.

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**spanning-tree bpdufilter** To enable bridge protocol data unit (BPDU) Filtering on the interface, use the spanning-tree bpdufilter command. To return to the default settings, use the no form of this command. `spanning-tree bpdufilter {enable	disable}`  `no spanning-tree bpdufilter` Syntax Description enable Enables BPDU Filtering on this interface. disable Disables BPDU Filtering on this interface. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 111.	**spanning-tree bpdufilter** The spanning-tree bpdufilter command controls bridge protocol data unit (BPDU) filtering on the configuration mode interface. BPDU filtering is disabled by default. Ports with BPDU filtering enabled drop inbound BPDUs and do not send BPDUs. Enabling BPDU filtering on a port not connected to a host can result in loops as the port continues forwarding data while ignoring inbound BPDU packets. • spanning-tree bpdufilter enabled enables BPDU filtering. • spanning-tree bpdufilter disabled disables BPDU filtering by removing the spanning-tree bpdufilter command from *running-config*. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 996.